



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
13.05.2020 Bulletin 2020/20

(51) Int Cl.:
H04L 12/18 (2006.01) **H04L 29/08 (2006.01)**
H04L 29/06 (2006.01)

(21) Application number: **19217248.4**

(22) Date of filing: **14.04.2016**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

• **GRIFFIN, Keith**
San Jose, CA California 95134 (US)

(30) Priority: **17.04.2015 US 201514690039**

(74) Representative: **Hirsz, Christopher Stanislaw et al**
Mathys & Squire LLP
The Shard
32 London Bridge Street
London SE1 9SG (GB)

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
16720239.9 / 3 284 211

Remarks:

This application was filed on 17-12-2019 as a divisional application to the application mentioned under INID code 62.

(71) Applicant: **Cisco Technology, Inc.**
San Jose, CA 95134 (US)

(72) Inventors:
• **ROSENBERG, Jonathan D.**
San Jose, CA California 95134 (US)

(54) **HANDLING CONFERENCES USING HIGHLY-DISTRIBUTED AGENTS**

(57) Systems, methods, and non-transitory computer-readable storage media for directing conferences to on-premises media agents and cloud-based media agents. A centralized controller and a plurality of cloud-based media agents are disposed within a cloud or other remote or virtual network, and a plurality of on-premises media agents are provided on-premises. The centralized controller manages the directing and forwarding of conferences and is in communication with the cloud-based media agents and the on-premises media agents.




FIG. 1

Description**TECHNICAL FIELD**

[0001] The present technology pertains to handling conferences by media agents.

BACKGROUND

[0002] Conferences, such as a teleconference having multiple callers requesting to join the teleconference, require a plurality of media agents, such as MCUs, to carry out a conference with multiple parties. Existing teleconferencing systems are either on-premises or in the cloud. It is difficult to anticipate the amount of resources needed for conferences, so companies are often forced to purchase the amount of on-premises or cloud capacity needed for peak capacity.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] In order to describe the manner in which the above-recited and other advantages and features of the disclosure can be obtained, a more particular description of the principles briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only exemplary embodiments of the disclosure and are not therefore to be considered to be limiting of its scope, the principles herein are described and explained with additional specificity and detail through the use of the accompanying drawings in which:

FIG. 1 illustrates an example block diagram according to some aspects of the subject technology;

FIG. 2 illustrates an example method embodiment;

FIG. 3 illustrates an example method embodiment;

FIG. 4 illustrates an example block diagram including a load balancer fronting each group of media agents, according to some aspects of the subject technology;

FIG. 5 illustrates an example method embodiment; and

FIGs. 6A and 6B illustrate example system embodiments according to some aspects of the subject technology.

DESCRIPTION OF EXAMPLE EMBODIMENTS

[0004] Various embodiments of the disclosure are discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and con-

figurations may be used without parting from the spirit and scope of the disclosure.

OVERVIEW

[0005] An elastic model for handling conferences includes a plurality of on-premises media agents and a plurality of cloud-based media agents that are used together to perform conferences. A "cloud" or other remote network includes a plurality of cloud-based media agents and a centralized controller. The centralized controller receives requests to join conferences. The controller can be cloud-based or can be based at any location as a standalone entity that is securely reachable over a network. The centralized controller provides a requester with instructions for attempting a first connection. This can be an attempt to connect to a media agent on-premises or cloud-based. If the first connection is not made, then a second connection can be attempted. The first connection may not be made due to several factors, such as the media agents meeting or exceeding capacity. A second connection attempt can then be made to an on-premises or cloud-based media agent. The controller can manage the capacity of the media agents, or a stun load balancer can be at the front end of each group of media agents to monitor the capacity of the media agents.

DESCRIPTION

[0006] A computer network is a geographically distributed collection of nodes interconnected by communication links and segments for transporting data between endpoints, such as personal computers and workstations. Many types of networks are available, with the types ranging from local area networks (LANs) and wide area networks (WANs) to overlay and software-defined networks, such as virtual extensible local area networks (VXLANS).

[0007] LANs typically connect nodes over dedicated private communications links located in the same general physical location, such as a building or campus. WANs, on the other hand, typically connect geographically dispersed nodes over long-distance communications links, such as common carrier telephone lines, optical light-paths, synchronous optical networks (SONET), or synchronous digital hierarchy (SDH) links. LANs and WANs can include layer 2 (L2) and/or layer 3 (L3) networks and devices.

[0008] The Internet is an example of a WAN that connects disparate networks throughout the world, providing global communication between nodes on various networks. The nodes typically communicate over the network by exchanging discrete frames or packets of data according to predefined protocols, such as the Transmission Control Protocol/Internet Protocol (TCP/IP). In this context, a protocol can refer to a set of rules defining how the nodes interact with each other. Computer networks may be further interconnected by an intermediate net-

work node, such as a router, to extend the effective "size" of each network.

[0009] Network virtualization allows hardware and software resources to be combined in a virtual network. For example, network virtualization can allow multiple numbers of VMs to be attached to the physical network via respective virtual LANs (VLANs). The VMs can be grouped according to their respective VLAN, and can communicate with other VMs as well as other devices on the internal or external network.

[0010] Cloud computing can also be provided in one or more networks to provide computing services using shared resources. Cloud computing can generally include Internet-based computing in which computing resources are dynamically provisioned and allocated to client or user computers or other devices on-demand, from a collection of resources available via the network (e.g., "the cloud"). Cloud computing resources, for example, can include any type of resource, such as computing, storage, and network devices, virtual machines (VMs), etc. For instance, resources may include service devices (firewalls, deep packet inspectors, traffic monitors, load balancers, etc.), compute/processing devices (servers, CPUs, memory, brute force processing capability), storage devices (e.g., network attached storages, storage area network devices), etc. In addition, such resources may be used to support virtual networks, virtual machines (VM), databases, applications (Apps), etc.

[0011] Cloud computing resources may include a "private cloud," a "public cloud," and/or a "hybrid cloud." A "hybrid cloud" can be a cloud infrastructure composed of two or more clouds that inter-operate or federate through technology. In essence, a hybrid cloud is an interaction between private and public clouds where a private cloud joins a public cloud and utilizes public cloud resources in a secure and scalable manner. Cloud computing resources can also be provisioned via virtual networks in an overlay network, such as a VXLAN.

[0012] The disclosed technology addresses the need in the art for an elastic model for handling conferences to provide adequate resources for average and peak bandwidth using a cloud-based controller and to provide conferencing capabilities to clients seeking to join a conference when media agents on-premises and/or media agents in the cloud are at capacity. Disclosed are systems, methods, and computer-readable storage media for clients of conferences to be handled by on-premises and cloud-based media agents. A cloud-based controller receives incoming requests to join conferences and selects an on-premises media agent or a cloud-based media agent to handle the request, depending upon a number of factors, such as resource bandwidth, call type, caller information and other call-related data. For example, all requests to join conferences may go to on-premises media agents until capacity is met or exceeded, and then the conferences may be handled by cloud-based media agents so the client can still join the conference even when the media agents on-premises have reached

capacity. Another example is for conferences of a first media type (e.g., audio conferences) to be handled by cloud-based media agents, and to direct conferences of a second media type (e.g., video conferences) to be handled by on-premises media agents. In still a further example, the controller has logic that determines which media agents are more secure, and accordingly directs secure conferences to the secure media agents.

[0013] A brief introductory description of exemplary systems and networks, as illustrated in FIG. 1, is disclosed herein. A detailed description of the various system flows, related concepts, and exemplary variations, will then follow. These variations shall be described herein as the various embodiments are set forth. The disclosure now turns to FIG. 1.

[0014] FIG. 1 illustrates an example block diagram according to some aspects of the subject technology. An on-premises enterprise 110 has a plurality of on-premises media agents 111, 112 and 113. Media agents perform media (e.g., audio, video and share) packet processing operations in support of conference sessions in which clients participate, such as, but not limited to, media packet mixing, switching, encoding, decoding and transcoding. Although only three media agents are shown and described for illustrative and descriptive purposes, it should be understood that any number of media agents can be employed to achieve the desired conferencing capabilities for a particular premise. Also, although only a single enterprise 110 is shown having one cluster of agents 111, 112, 113, the techniques are likewise applicable to any number of enterprises, each having any number of media agents or clusters of media agents. In accordance with the techniques herein, enterprises having multiple locations that each have a plurality of media agents and are geographically distributed can employ the conferencing system herein to control overall enterprise conferencing capabilities for all locations using a cloud-based controller. Audio, video and sharing capabilities are handled using the cloud-based controller.

[0015] In accordance with some aspects of the subject technology, an enterprise (e.g., a company) can purchase sufficient capacity for average demand on-premises, instead of peak demand, and purchase capacity in the cloud for peak capacity, to be used on an as-needed basis. In this manner, a fixed capacity of licensed computing is deployed on-premises at the enterprise, and the remaining capacity can be provided by the cloud resources.

[0016] A conference session (also referred to more generally as a "conference") may involve the exchange of one or more of audio, video and sharing of resources (such as screen or application share) between any number of clients as well as the sharing of content by one participant to one or more other clients. Such shared content may include documents, presentations, audio content, video content, etc. A conference sessions including screen and/or application share can be handled by the cloud-based controller in accordance with the

teachings herein. The cloud controller can thus handle conferences with audio, video and sharing capabilities. The sharing capabilities can be encoded and handled as video in at least some embodiments.

[0017] A controller 130 resides in the "cloud" 125 which can be an appropriate data center. To control the system, the controller 130 performs control plane signal operations/functions using such features and functions afforded by, for example but not limited to, the Session Initiation Protocol (SIP), H323 access, rosters, and conference control options (including, e.g., mute, kick, etc.). The controller also performs orchestration of the conferences by controlling the connections between on-premises media agents and cloud-based media agents to ensure that a particular conference is fully connected and provides the necessary audio and video processing functions. The controller can have logic that allows it to select on-premises media agents or cloud-based media agents in accordance with some aspects of the subject technology.

[0018] The cloud 125 also includes a plurality of media agents 121, 122 and 123 that are cloud-based and reside within the cloud 125 as opposed to on-premises. Although only three media agents are shown and described for illustrative and descriptive purposes, it should be understood that any number of media agents can be employed to achieve the desired conferencing capabilities for a particular premise through the use of cloud-based media agents in addition to on-premises media agents. In this manner, an enterprise can utilize the on-premises media agents and the cloud-based media agents to accommodate their conferencing needs without requiring an enterprise to maintain sufficient media agents on-premises to satisfy peak conferencing needs, as the cloud-based media agents can be used on an as-needed basis, or depending upon particular factors that the enterprise has determined to control the media agent that handle the conferences.

[0019] The system of FIG. 1 provides a dynamic, elastic technique to allow for maximum utilization of collaboration in resources between on-premises and the cloud. As will be described in detail below, techniques presented herein dynamically configure highly-distributed resources in a conference system, including the cloud-based controller, the on-premises media agents and the cloud-based media agents to support conference sessions initiated by clients and direct conferences to the appropriate on-premises or cloud-based media agents. The resources are referred to as "highly-distributed" to describe geographically distributed resources.

[0020] In support of these techniques, the controller 130, on-premises media agents 111, 112, 113 and cloud-based media agents 121, 122, 123 implement and interact with each other using a variety of communication protocols to establish conference sessions and exchange media streams including packets of data in conference sessions. Such communication protocols include, but are not limited to: the Interactive Connectivity Establishment (ICE) protocol; the Session Traversal Utilities for Network

Address Translation/Translator (NAT) (STUN) protocol modified or extended to use STUN URLs, in accordance with techniques presented herein; the User Datagram Protocol (UDP); and the Real-Time Transport protocol (RTP). The techniques described herein use the aforementioned protocols by way of example, only; other similar protocols may also be used instead of or in combination with the mentioned protocols, as would be appreciated by one of ordinary skill in the relevant arts having access to the description presented herein.

[0021] The on-premises media agents 111, 112, 113 and cloud-based media agents 121, 122, 123 may be any number of media agents, up to the tens of thousands, and can be distributed geographically around the world. Similarly, clients seeking to join a conference may be located anywhere in the world. As such, the conference system is referred to as "highly distributed." A challenge to such a distributed arrangement is to construct a "best" media topology for assigning clients to conferences using on-premises media agents or cloud-based media agents.

[0022] Techniques presented herein address the challenges of assigning clients to conferences based upon a number of factors. For example, an enterprise may wish to control the use of media agents based on the type of data being exchanged in the conference. The enterprise may desire to have all audio conferences handled by on-premises media agents while all video conferences are handled by cloud-based media agents, or vice versa, so that the appropriate media agents can handle the various conferences. As another example, the enterprise may wish to simply set up their conferencing capabilities based on capacity of the media agents so that the on-premises media agents are used first until there are no longer any media agents available on-premises, and then the cloud-based media agents can be used when the media agents on-premises have reached capacity. Yet another example is to employ on-premises media agents for on-premises clients, and to employ cloud-based media agents for clients that are not on-premises so that the on-premises clients use the on-premises media agents and remote clients use the cloud-based media agents. And yet another example is handling conferences based on the security of the conference, so that secure conferences are handled by on-premises media agents, and non-secure conferences are handled by cloud-based media agents, or vice-versa with the non-secure conferences handled by on-premises media agents and secure conferences handled by cloud-based media agents. There are several decisions in the controller logic to control routing of conferences to on-premises media agents and cloud-based media agents. The controller logic can determine which media agents are more secure and direct conferences accordingly. For example, if the on-premises media agents are more secure, the secure conferences are directed to the on-premises media agents, and if the cloud-based media agents are more secure, the secure conferences are directed to cloud-based media agents. The factors for determining how conferences

are directed to cloud-based media agents or on-premises media agents can be chosen by the enterprise or pre-selected by the manufacturer, or may be available for on-demand changes depending upon the needs and desired functionality in conferencing of the enterprise.

[0023] The advantages of a conferencing system that includes both on-premises media agents and cloud-based media agents include limitless conferencing capabilities, meaning there is no limit on the number of participants, and the conferencing system is able to satisfy the needs of traditional on-premises conferencing and cloud-based conferencing individually, as well as combined, by combining the two conferencing mechanisms into a single conferencing system.

[0024] FIG. 2 illustrates an example method embodiment. The procedure of FIG. 2 commences by 210 registering a plurality of on-premises media agents with a cloud-based controller. The plurality of on-premises media agents can, for example, be the media agents 111, 112, 113 shown in FIG. 1, and can include a single on-premises media agent up to any number of media agents residing on-premises at an enterprise. At 214 a plurality of cloud-based media agents are registered with the cloud-based controller. At 216, incoming requests to join a teleconference, which can be referred to generally as "calls", can be any request to join a conference, such as an audio or video teleconference, are handled by at least one on-premises media agent or at least one cloud-based media agent based upon a plurality of factors. The factors for controlling the handling of incoming requests to join a conference can include capacity for more clients at the on-premises media agents, security of the conference, media type of the conference, and other factors described herein and known in the relevant art as pertaining to conferencing sessions and conferencing systems. The conferences can be general teleconferences, such as audio or video conferences, and include multi-media conferences.

[0025] The handling of conferences by cloud-based media agents and on-premises media agents can be seamless to the clients. There are two approaches for directing conferences, a first in which the controller manages the switch. A second is the switch is handled by the client and a load balancer.

[0026] FIG. 3 illustrates an example method embodiment for handling conferences in accordance with some aspects of the subject technology where the controller manages the switching of conferences. In this embodiment, conferences are handled by media agents that are located on-premises until they have reached maximum capacity, and then the cloud-based controller handles the conferences or reallocates some data to cloud-based media agents. The capacity threshold can indicate that the media agents cannot handle further audio or video streams. If the capacity threshold is met, then the centralized controller can locate cloud-based media agents to handle media requests from new clients requesting to join the conference.

[0027] As described herein, the threshold capacity is only one factor that can be used in determining the media agent that handles a requested conference. The handling can also occur based on the type of security of the conference, such that the secure conferences are handled by media agents that are deemed to be the most secure. For example, the controller logic can determine the media agents that are the most secure, and direct the secure calls to the most secure media agents. The threshold capacity is used as one example factor for determining the handling of conferences in the example embodiment of FIG. 3.

[0028] At 310, when a client device desires to join a conference, the client device transmits a message to a centralized controller requesting to join a teleconference (including an audio conference or a video conference). The centralized controller is cloud-based and resides in the cloud in the example embodiment. At 312, the centralized controller returns a media URL that can be used to identify an appropriate media agent for the client device to connect to the conference. The media agent identified in the media URL may be a cloud-based media agent or an on-premises media agent, depending upon the logic of the cloud-based controller for handling conferences by the appropriate media agents. The controller can optionally further command appropriate media agents to communicate with each other to facilitate the conference at 314. At 316, the media agents send a notification to the controller when they have reached a threshold capacity. At 318 it is determined by the cloud-based controller if all local media agents have reached their capacity threshold. When all on-premises media agents have not reached their threshold capacity at 318, the centralized controller returns to 312 and continues to provide media URLs to clients to be handled by on-premises media agents. When all on-premises media agents have reached their threshold capacity at 318, the centralized controller reallocates some data communications or directs future conference requests to cloud-based media agents at 320.

[0029] FIG. 4 illustrates an example block diagram including a STUN [Simple Traversal of UDP (Universal Datagram Protocol) through NATs (Network Address Translators)] load balancer (also referred to as a "load balancing server") fronting each group of media agents, according to some aspects of the subject technology. Although depicted as a STUN load balancer, it should be understood that any load balancer can be implemented in accordance with the teachings of the subject technology. Further, it is also contemplated that in alternate embodiments, the load balancer functionality can be absorbed into the media agent. Accordingly, when the load balancer receives the join request, it interrogates a shared database amongst the load balancers, and determines which if any media agent is already servicing the conference. If there is already one assigned, the load balancer redirects the client to that media agent. If not, the load balancer redirects the client to a media agent with avail-

able capacity.

[0030] The techniques presented herein dynamically configure highly-distributed resources to support conference sessions initiated by callers on an as-needed basis. In support of these techniques, the controller, media agents, callers and load balancers, implement and interact with each other using a variety of communication protocols to establish conference sessions and exchange media streams containing packets of data in the conference sessions. Such communication protocols include, but are not limited to: the Interactive Connectivity Establishment (ICE) protocol; the STUN protocol modified/extended to use STUN URLs, in accordance with techniques presented herein; the User Datagram Protocol (UDP) and the Real-Time Transport Protocol (RTP). The techniques described herein use the aforementioned protocols by way of example only; other similar protocols may also be used instead of or in combination with the mentioned protocols, as would be appreciated by one of ordinary skill in the relevant arts having access to the description presented herein.

[0031] FIG. 5 illustrates an example method embodiment for directing conferences using, for example, the conference system including STUN load balancers as shown in FIG. 4. At 510, when a client connects to the controller, the client is given a single STUN URL 510. When the client resolves this DNS from within the boundaries of the enterprise, it results in two distinct records at 512. The first record is a VIP (Virtual IP address) that points to a cluster of cloud-based agents. The other record, at a higher priority, points to a cluster of media agents on-premises. Fronting each of these clusters is a STUN load balancer as shown, however it should be understood that any load balancer can be implemented and/or that the load balancing functionality can be absorbed into the media agent itself.

[0032] At 514 the client uses the STUN to connect with higher priority to the on-premises media agent. At 516, the STUN load balancer monitors the load of the cloud-based media agents. At 518, the STUN load balancer determines if there is sufficient capacity for the on-premises media agents. When there is sufficient capacity, the request is handled by an on-premises media agent at 520 and then the user is connected to an on-premises media agent at 522. When there is not sufficient capacity at 518, the STUN load balancer can reject the STUN request at 524, with an error response which indicates a temporary capacity failure. The client uses this as a signal to instead try the second IP address at 526. The client device performs STUN to the second IP address at 528, which is a cloud-based media agent. Since there is enough capacity in the cloud, the load balancer can direct the request to an available media agent and the STUN transaction completes successfully. At 530 the client is connected to a cloud-based media agent.

[0033] The benefit of this example embodiment shown in FIG. 5 is that it does not require the controller to know the available capacity of the media agents and thus

scales to a larger number of clusters worldwide. This technique also works particularly well with split-horizon DNS where the STUN URL would return two VIPs for resolutions performed within the enterprise boundary, but only one (the cloud VIP) when the client is outside of the boundary of the enterprise. This model is primarily concerned with implementing cloud-based resources to account for on-premises media agents exceeding capacity. It is also possible for an organization to utilize available on-premises resources from a cloud-based meeting.

[0034] In another example embodiment, the handling of conferences by cloud-based media agents and on-premises media agents may be based on media modality rather than exceeding capacity of the media agents. For example, an on-premises meeting may elect to direct conferences of a specific media type or combination of modalities to the cloud while maintaining other resources on-premises. The cloud controller can be provided with logic to determine the best mode of directing conferences based upon the information associated with a particular conference. For example, if an audio conference is requested, the audio conference is handled by a first group of media agents, if a video conference is requested, the video conference is handled by a second group of media agents, and if a sharing conference is requested, the sharing conference is handled by a third group of media agents.

[0035] FIG. 6A, and FIG. 6B illustrate example system embodiments according to some aspects of the subject technology. The more appropriate embodiment will be apparent to those of ordinary skill in the art when practicing the present technology. Persons of ordinary skill in the art will also readily appreciate that other system embodiments are possible.

[0036] FIG. 6B illustrates a conventional system bus computing system architecture 600 wherein the components of the system are in electrical communication with each other using a bus 605. Exemplary system 600 includes a processing unit (CPU or processor) 610 and a system bus 605 that couples various system components including the system memory 615, such as read only memory (ROM) 620 and random access memory (RAM) 625, to the processor 610. The system 600 can include a cache of high-speed memory connected directly with, in close proximity to, or integrated as part of the processor 610. The system 600 can copy data from the memory 615 and/or the storage device 630 to the cache 612 for quick access by the processor 610. In this way, the cache can provide a performance boost that avoids processor 610 delays while waiting for data. These and other modules can control or be configured to control the processor 610 to perform various actions. Other system memory 615 may be available for use as well. The memory 615 can include multiple different types of memory with different performance characteristics. The processor 610 can include any general purpose processor and a hardware module or software module, such as module 1 632,

module 2 634, and module 3 636 stored in storage device 630, configured to control the processor 610 as well as a special-purpose processor where software instructions are incorporated into the actual processor design. The processor 610 may essentially be a completely self-contained computing system, containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0037] To enable user interaction with the computing device 600, an input device 645 can represent any number of input mechanisms, such as a microphone for speech, a touch-sensitive screen for gesture or graphical input, keyboard, mouse, motion input, speech and so forth. An output device 635 can also be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems can enable a user to provide multiple types of input to communicate with the computing device 600. The communications interface 640 can generally govern and manage the user input and system output. There is no restriction on operating on any particular hardware arrangement and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[0038] Storage device 630 is a non-volatile memory and can be a hard disk or other types of computer readable media which can store data that are accessible by a computer, such as magnetic cassettes, flash memory cards, solid state memory devices, digital versatile disks, cartridges, random access memories (RAMs) 625, read only memory (ROM) 620, and hybrids thereof.

[0039] The storage device 630 can include software modules 632, 634, 636 for controlling the processor 610. Other hardware or software modules are contemplated. The storage device 630 can be connected to the system bus 605. In one aspect, a hardware module that performs a particular function can include the software component stored in a computer-readable medium in connection with the necessary hardware components, such as the processor 610, bus 605, display 635, and so forth, to carry out the function.

[0040] FIG. 6A illustrates a computer system 650 having a chipset architecture that can be used in executing the described method and generating and displaying a graphical user interface (GUI). Computer system 650 is an example of computer hardware, software, and firmware that can be used to implement the disclosed technology. System 650 can include a processor 655, representative of any number of physically and/or logically distinct resources capable of executing software, firmware, and hardware configured to perform identified computations. Processor 655 can communicate with a chipset 660 that can control input to and output from processor 655. In this example, chipset 660 outputs information to output 665, such as a display, and can read and write information to storage device 670, which can include magnetic media, and solid state media, for example. Chipset 660 can also read data from and write data

to RAM 675. A bridge 680 for interfacing with a variety of user interface components 685 can be provided for interfacing with chipset 660. Such user interface components 685 can include a keyboard, a microphone, touch detection and processing circuitry, a pointing device, such as a mouse, and so on. In general, inputs to system 650 can come from any of a variety of sources, machine generated and/or human generated.

[0041] Chipset 660 can also interface with one or more communication interfaces 690 that can have different physical interfaces. Such communication interfaces can include interfaces for wired and wireless local area networks, for broadband wireless networks, as well as personal area networks. Some applications of the methods for generating, displaying, and using the GUI disclosed herein can include receiving ordered datasets over the physical interface or be generated by the machine itself by processor 655 analyzing data stored in storage 670 or 675. Further, the machine can receive inputs from a user via user interface components 685 and execute appropriate functions, such as browsing functions by interpreting these inputs using processor 655.

[0042] It can be appreciated that exemplary systems 600 and 650 can have more than one processor 610 or be part of a group or cluster of computing devices networked together to provide greater processing capability.

[0043] As one of ordinary skill in the art will readily recognize, the examples and technologies provided above are simply for clarity and explanation purposes, and can include many additional concepts and variations.

[0044] For clarity of explanation, in some instances the present technology may be presented as including individual functional blocks including functional blocks comprising devices, device components, steps or routines in a method embodied in software, or combinations of hardware and software.

[0045] In some embodiments the computer-readable storage devices, mediums, and memories can include a cable or wireless signal containing a bit stream and the like. However, when mentioned, non-transitory computer-readable storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[0046] Methods according to the above-described examples can be implemented using computer-executable instructions that are stored or otherwise available from computer readable media. Such instructions can comprise, for example, instructions and data which cause or otherwise configure a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. Portions of computer resources used can be accessible over a network. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, firmware, or source code. Examples of computer-readable media that may be used to store instructions, information used, and/or information created during methods according to de-

scribed examples include magnetic or optical disks, flash memory, USB devices provided with non-volatile memory, networked storage devices, and so on.

[0047] Devices implementing methods according to these disclosures can comprise hardware, firmware and/or software, and can take any of a variety of form factors. Typical examples of such form factors include laptops, smart phones, small form factor personal computers, personal digital assistants, rackmount devices, standalone devices, and so on. Functionality described herein also can be embodied in peripherals or add-in cards. Such functionality can also be implemented on a circuit board among different chips or different processes executing in a single device, by way of further example.

[0048] The instructions, media for conveying such instructions, computing resources for executing them, and other structures for supporting such computing resources are means for providing the functions described in these disclosures.

[0049] Although a variety of examples and other information was used to explain aspects within the scope of the appended claims, no limitation of the claims should be implied based on particular features or arrangements in such examples, as one of ordinary skill would be able to use these examples to derive a wide variety of implementations. Further and although some subject matter may have been described in language specific to examples of structural features and/or method steps, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to these described features or acts. For example, such functionality can be distributed differently or performed in components other than those identified herein. Rather, the described features and steps are disclosed as examples of components of systems and methods within the scope of the appended claims. Moreover, claim language reciting "at least one of" a set indicates that one member of the set or multiple members of the set satisfy the claim.

[0050] Therefore, the scope of the invention is to be limited only by the following claims. However, embodiments are also described in the following numbered clauses:

1. A method comprising:

registering a plurality of on-premises media agents with a cloud-based controller, the plurality of on-premises media agents disposed at a physical location of an enterprise and the cloud-based controller disposed at a cloud location that is remote from the physical location of the enterprise;
 registering a plurality of cloud-based media agents with the cloud-based controller, the plurality of cloud-based media agents disposed at the cloud location that is remote from the physical location of the enterprise;
 receiving a request from a client device to join

a conference at the cloud-based controller; and selecting at least one on-premises media agent or at least one cloud-based media agent to handle the request depending upon at least one factor associated with the conference that the client device is requesting to join.

2. The method of clause 1, wherein directing the request comprises determining that a capacity of on-premises media agents exceeds a limit.

3. The method of clause 1, wherein directing the request comprises determining that a first connection attempt to a first media agent cannot be made and connecting to a second media agent.

4. The method of clause 1, wherein the factor associated with the conference comprises a capacity of the media agents.

5. The method of clause 1, wherein the factor associated with the conference comprises a security type of the conference, such that secure conferences are directed to on-premises media agents when the on-premises media agents are deemed to be more secure than the cloud-based media agents.

6. The method of clause 1, wherein the factor associated with the conference comprises a security type of the conference, such that secure conference are directed to cloud-based media agents when the cloud-based media agents are deemed to be more secure than the on-premises media agents.

7. The method of clause 1, wherein the factor associated with the conference comprises a media type of the conference, such that conferences of a first media type are directed to on-premises media agents and conferences of a second media type are directed to cloud-based media agents.

8. A system comprising:

a plurality of on-premises media agents disposed at a physical location of an enterprise;
 a plurality of cloud-based media agents disposed at a cloud location that is remote from the physical location of an enterprise; and
 a controller disposed at the cloud location that selects at least one on-premises media agent or at least one cloud-based media agent to handle a request to join a conference depending upon certain factor associated with the conference.

9. The system of clause 8, wherein selecting is based upon capacity of on-premises media agents exceeding a limit.

10. The system of clause 8, further comprising an STUN load balancer that monitors a capacity of each of the plurality of on-premises media agents and a cloud-based STUN load balancer that monitors a capacity of each of the plurality of cloud-based media agents.

11. The system of clause 8, wherein the conference is directed based upon a media type or a security type of the conference.

12. The system of clause 8 wherein secure conferences are directed to at least one of the plurality of on-premises media agents and non-secure conferences are directed to at least one of the plurality of cloud-based media agents.

13. A method comprising:

receiving, by a centralized controller, a message from a requester requesting to join a conference; attempting a first connection to a first media agent, the first media agent comprising at least one of a plurality of on-premises media agents at an enterprise; and attempting a second connection to a second media agent when a result of attempting the first connection to the first media agent is not successful, the second media agent comprising at least one of a plurality of cloud-based media agents at a cloud location remote from the enterprise.

14. The method of clause 13, wherein the result of attempting the first connection to the first media agent is not successful when the first media agent has reached or exceeded a capacity limit.

15. The method of clause 13, wherein a capacity of the cloud-based media agents and the on-premises media agents is monitored by one or more of: the controller, a cloud-based STUN load balancer, and an on-premises STUN load balancer.

Claims

1. A method comprising:

sending from a client device a request to join a conference at a cloud-based controller; storing, in response to a response from the controller, (a) a first address of a cluster of on-premises media agents disposed at a physical location of an enterprise, and (b) a second address of cloud-based media agents disposed at a cloud location that is remote from the physical location of the enterprise; selecting at least one on-premises media agent or at least one cloud-based media agent to handle the request, comprising:

selecting the first address in response to the cluster of on-premises media agents having sufficient capacity to handle the request; and selecting the second address in response to the cluster of on-premises media agents

lacking sufficient capacity to handle the request.

2. The method of claim 1, further comprising receiving, in response to the sending, a Simple Traversal of UDP (Universal Datagram Protocol) through NATs (Network Address Translators) (STUN) address of a STUN load balancer.

3. The method of either claim 1 or claim 2, further comprising, in response to the receiving, the first and second address.

4. The method of claim 3, further comprising setting the first address to a higher priority than the second address.

5. The method of any preceding claim, wherein the selecting at least one on-premises media agent or at least one cloud-based media agent to handle the request comprises determining whether the cluster of on-premises media agents has enough capacity to handle the request.

6. The method of claim 5, wherein the selecting the first address is in response to at least a positive result of the determining.

7. The method of either claim 5 or claim 6, wherein the selecting the second address is in response to at least a negative result of the determining.

8. A device comprising:

a memory storing instructions; a processor programmed to execute the instructions to perform the steps of a method according to any one of the previous claims.

9. A computer readable medium including executable instructions which, when executed in a processing system, cause the processing system to perform all the steps of a method according to any one of claims 1 to 7.




FIG. 1

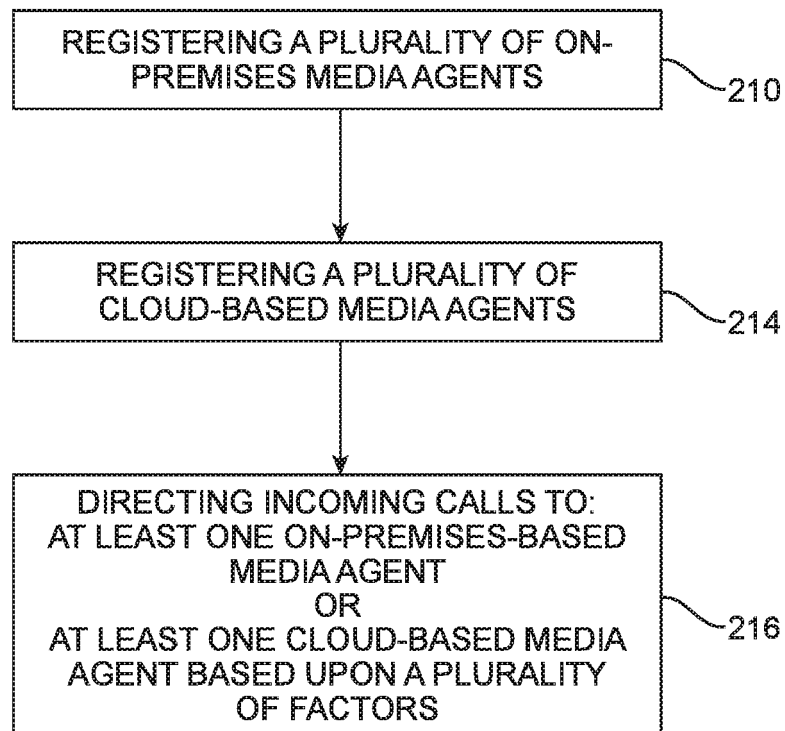


FIG. 2

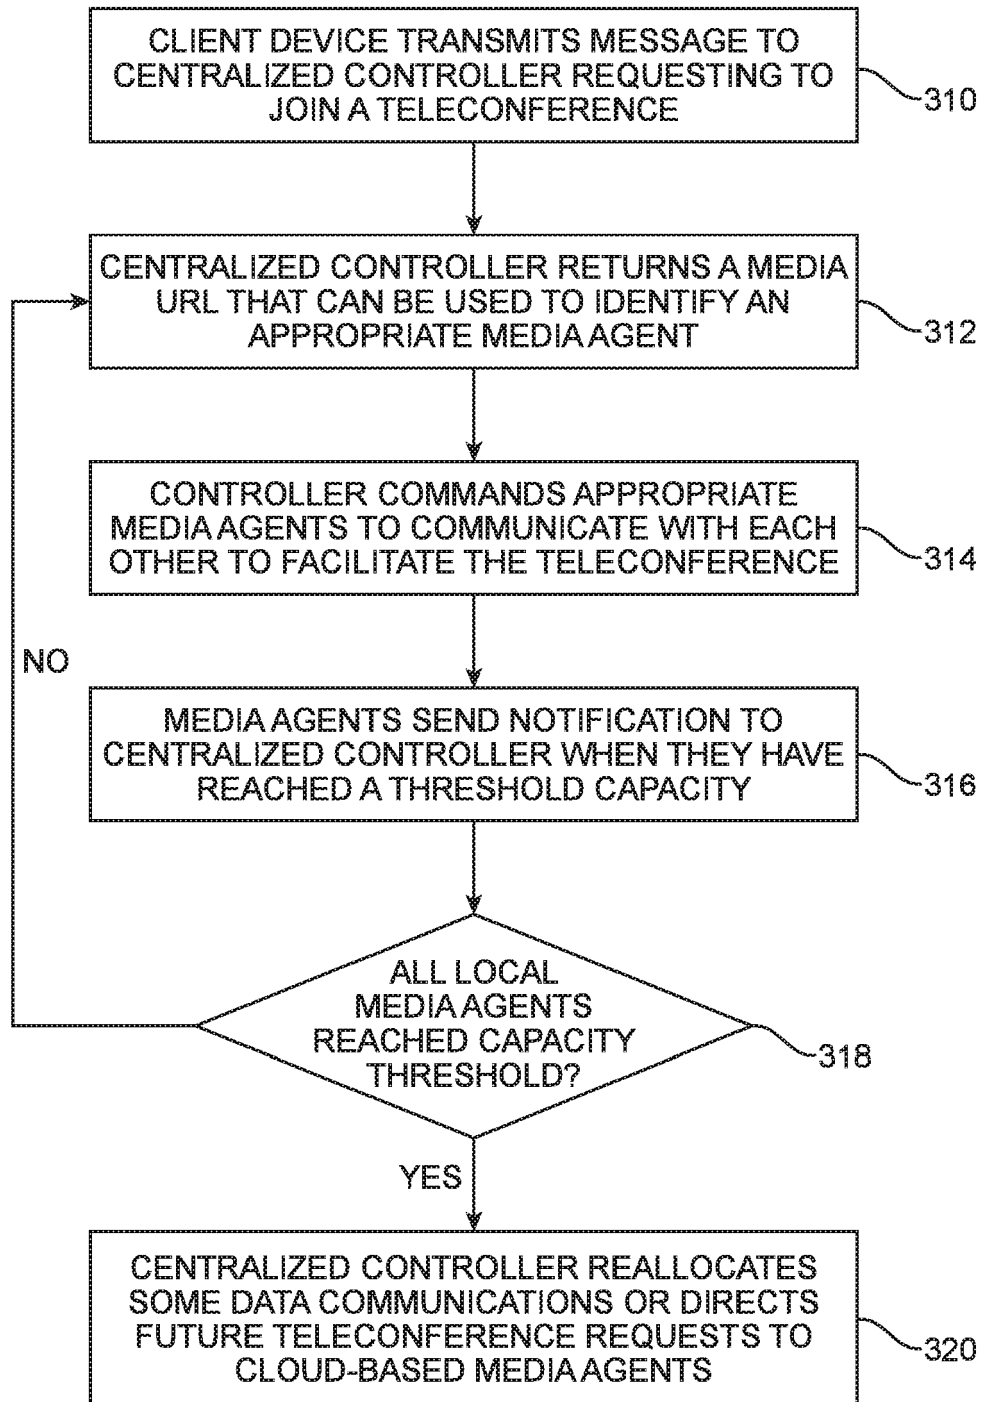


FIG. 3

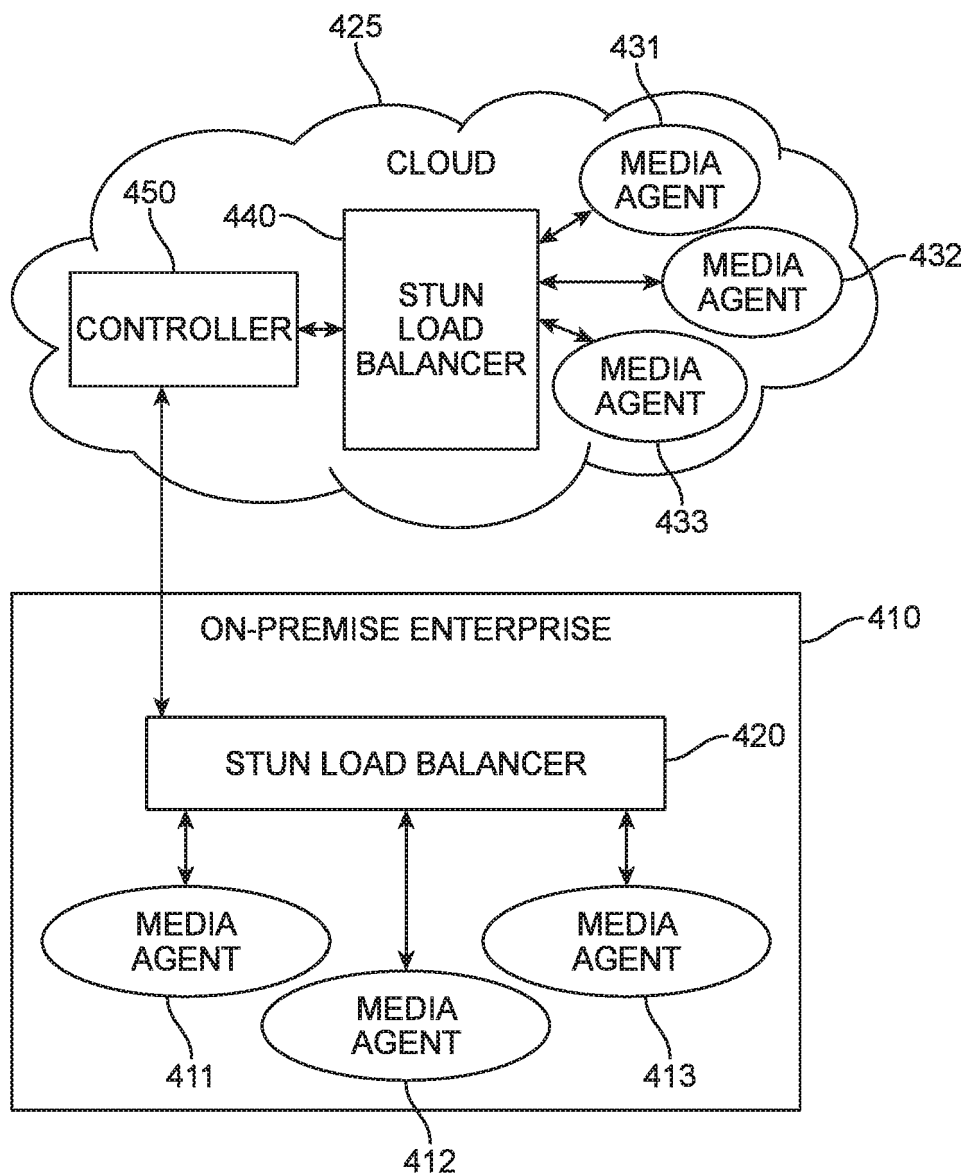


FIG. 4

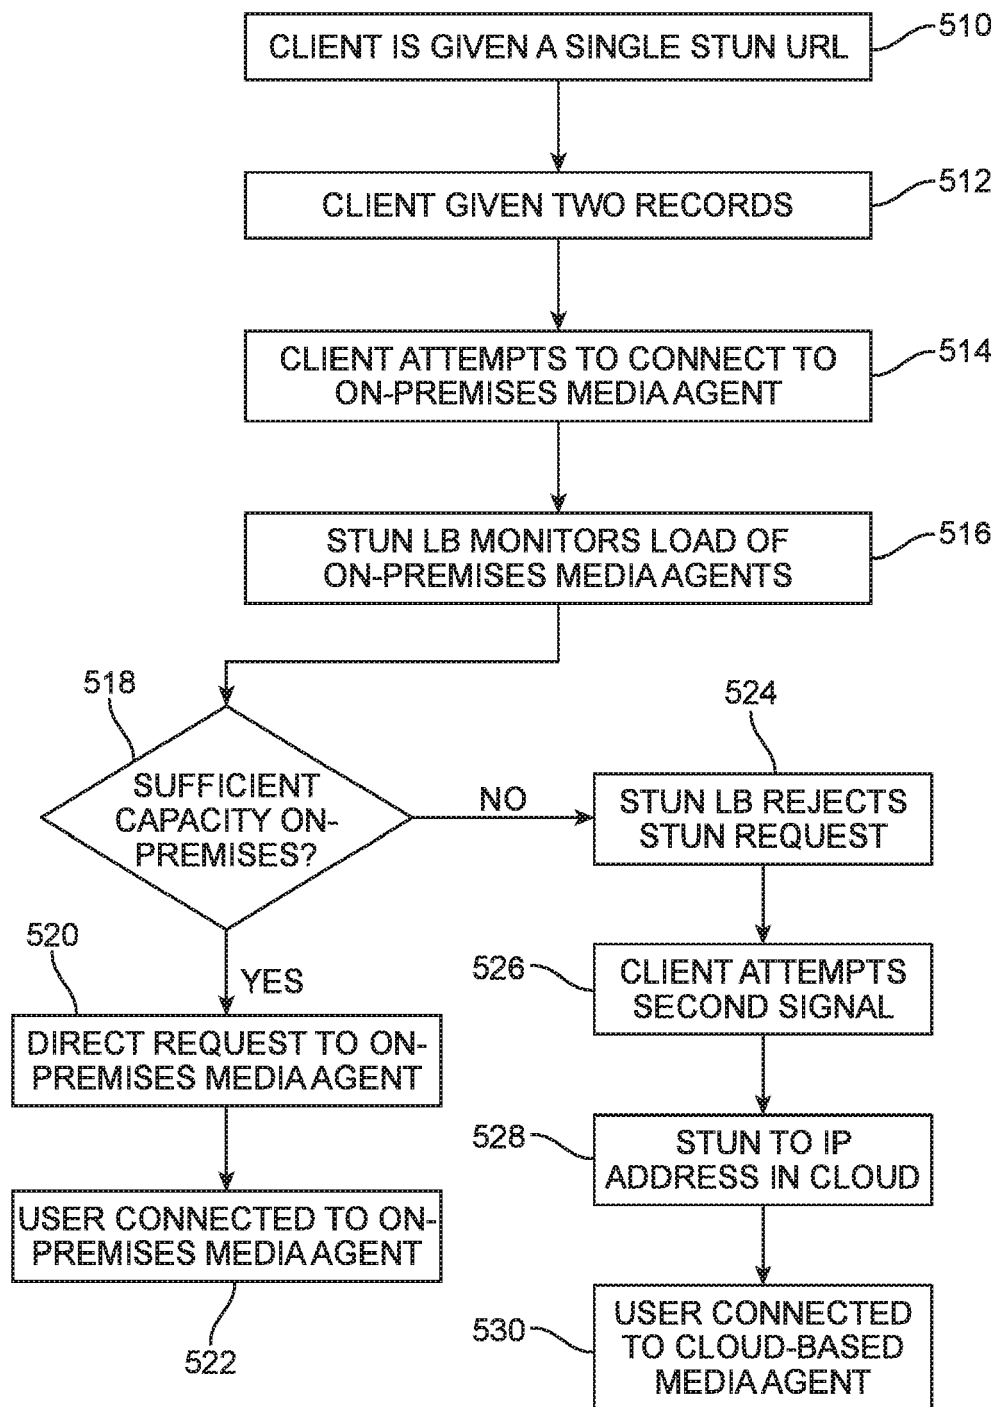


FIG. 5




FIG. 6A




FIG. 6B



EUROPEAN SEARCH REPORT

Application Number
EP 19 21 7248

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 6 646 997 B1 (BAXLEY WARREN E [US] ET AL) 11 November 2003 (2003-11-11) * column 2, line 45 - column 11, line 16; figures 1-2 *	1-9	INV. H04L12/18 H04L29/08 H04L29/06
A	WO 2014/118736 A2 (TAWQK CORP LTD [GB]) 7 August 2014 (2014-08-07) * paragraph [00123] - paragraph [00132]; figures 1,12,13 * * paragraph [0039] *	1-9	
A	US 2014/157338 A1 (PEARCE CHRISTOPHER E [US]) 5 June 2014 (2014-06-05) * paragraph [0018] - paragraph [0039]; figures 1-6 *	1-9	
A,P	US 2015/358472 A1 (ROSENBERG JONATHAN D [US]) 10 December 2015 (2015-12-10) * paragraph [0020] - paragraph [0028] * * paragraph [0102] - paragraph [0116]; figure 14 *	1-9	
			TECHNICAL FIELDS SEARCHED (IPC)
			H04L
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 20 February 2020	Examiner Fischer, Erik
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 19 21 7248

5

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

20-02-2020

10

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6646997 B1	11-11-2003	US 6646997 B1	11-11-2003
		US 2004085914 A1	06-05-2004
-----	-----	-----	-----
WO 2014118736 A2	07-08-2014	US 2015381666 A1	31-12-2015
		WO 2014118736 A2	07-08-2014
-----	-----	-----	-----
US 2014157338 A1	05-06-2014	NONE	
-----	-----	-----	-----
US 2015358472 A1	10-12-2015	NONE	
-----	-----	-----	-----

15

20

25

30

35

40

45

50

55

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82